An Expressively Complete Linear Time Temporal Logic for Mazurkiewicz Traces[1]

P. S. Thiagarajan
SPIC Mathematical Institute
92 G.N.Chetty Road Madras 600017
India
pst@smi.ernet.in

I. Walukiewicz[2]
Institute of Informatics, Warsaw University
Banacha 2, 02-097 Warsaw
Poland
igw@mimuw.edu.pl

Abstract

A basic result concerning LTL, the propositional temporal logic of linear time, is that it is expressively complete; it is equal in expressive power to the first order theory of sequences. We present here a smooth extension of this result to the class of partial orders known as Mazurkiewicz traces. These partial orders arise in a variety of contexts in concurrency theory and they provide the conceptual basis for many of the partial order reduction methods that have been developed in connection with LTL-specifications.

We show that LTrL, our linear time temporal logic, is equal in expressive power to the first order theory of traces when interpreted over (finite and) infinite *traces. This result fills a prominent gap in the existing logical theory of infinite traces. LTrL also provides a syntactic characterisation of the so called trace consistent (robust) LTL-specifications. These are specifications expressed as LTL formulas that do not distinguish between different linearisations of the same trace and hence are amenable to partial order reduction methods.*

1 Introduction

A basic result concerning LTL, the propositional temporal logic of linear time, is that it is expressively complete; it is equal in expressive power to the first order theory of sequences [10, 6, 23]. Here we present a natural extension of this result to the class of labelled partial orders known as Mazurkiewicz traces.

To motivate this extension, we first note that as suggested by Pnueli [16], LTL is often interpreted over the runs of a distributed system. It is well known that these runs can be grouped together into equivalence classes using the causal independence of actions performed by different agents at separate locations; two runs are equated just in case they constitute two different linearisations of the same partially ordered stretch of behaviour. Thus each equivalence class corresponds to – all possible – linearisations of a partial order. In many settings, the partial orders that arise in this fashion are Mazurkiewicz traces. It is also often the case that the property expressed by an LTL formula is insensitive to the choice of linearisations in the sense that either all members of an equivalence class of runs satisfy the formula or none do. For verifying such properties it suffices to check that it holds for just one member of each equivalence class. The resulting savings in the verification process can be often substantial. This is the insight that underlies many of the so-called partial order reduction techniques [9, 14, 22].

There is an alternative way to exploit the non-sequential nature of the behaviour of distributed systems and the resulting partial order based reduction methods. It consists of developing temporal logics that can be directly interpreted over partial orders corresponding to the equivalence classes of runs of a distributed system. This is one of the main motivations for studying linear time temporal logics that are interpreted over traces. Yet another motivation is that traces are intimately related to basic objects in concurrency theory such as Petri nets and event structures [13].

Starting with [18] a number of such logics have been proposed in the literature [1, 3, 12, 17]. The study of these logics has so far left open two important expressiveness issues, one theoretical and the other prag-

[1]This work was done at BRICS, Basic Research in Computer Science, Centre of the Danish National Research Foundation, Computer Science Department, Aarhus University, Denmark.
[2]The author was partially supported by Polish KBN grant No. 8 T11C 002 11.

matic. From a theoretical standpoint, it has not been possible so far to exhibit a temporal logic patterned after LTL that has the same expressive power as the first order theory of infinite traces. This has been been an annoying gap in an otherwise smooth generalisation of the theory of ω-sequences to the theory of infinite traces [8, 4]. From a pragmatic standpoint, the "local" nature of the semantics of the logics that have been proposed so far makes it impossible to express arbitrary global liveness and safety properties. (More material on these issues can be found in [11].)

The temporal logic that we propose here, denoted $LTrL$, fulfils both these criteria. It is expressively complete for both finite and infinite traces. And we can transparently formulate within $LTrL$, global liveness and safety properties of all kinds. The existence of such a temporal logic is not guaranteed in advance since the class of partial orders as a whole does not admit an expressively complete temporal logic [7]. The only comparable previous result concerning traces is due to Ebinger [3]. His logic, denoted TLPO, has both previous state and Since modalities. These past modalities are used extensively in the proof of the fact that TLPO is expressively complete when interpreted over *finite* traces. This proof does not extend to infinite traces. In contrast, $LTrL$ uses just a very restricted past state modality and as stated already, it is expressively complete over the domain of infinite traces.

One consequence of our main result is that $LTrL$ captures exactly the so called trace consistent (robust) LTL-definable properties [19, 15]. These are properties that are naturally amenable to partial order reduction methods. A second consequence is that the satisfiability problem for $LTrL$ is decidable. The decision procedure we obtain is non-elementary.

At present, we do not have an elementary decision procedure for $LTrL$ and this is a limitation shared by TLPO. Proving the existence of such a procedure seems to be a challenging problem and its outcome will strongly influence the applicability of our logic as a specification formalism.

In the next section we introduce traces. The first order theory of traces as well as the syntax and semantics of $LTrL$ are presented in section 3. This leads to the formulation of the main result and its corollaries. The major ingredients of the proof of the main result are: An observation in [4] that makes crucial use of [6], a decomposition result for infinite traces, an easy version of the Feferman-Vaught theorem for generalised products [5] and a new normal form linearisation of traces. Indeed, it is the use of the Feferman-Vaught result and the new normal form that takes us past the key technical hurdles. The proof is presented in section 4.

2 Traces

A (Mazurkiewicz) *trace alphabet* is a pair (Σ, I) where Σ is a finite set of actions and $I \subseteq \Sigma \times \Sigma$ is an irreflexive and symmetric *independence relation.* $D = (\Sigma \times \Sigma) - I$ is called the dependency relation. Through the rest of the paper we fix a trace alphabet (Σ, I) and we will often refer to it implicitly. We let a, b range over Σ.

We shall view (Mazurkiewicz) trace as a restricted Σ-labelled poset. Let $(E, \leq, \lambda)$ be a Σ-labelled poset. In other words, $(E, \leq)$ is a poset and $\lambda : E \to \Sigma$ is a labelling function. For $Y \subseteq E$ we define $\downarrow Y = \{x \mid \exists y \in Y.\, x \leq y\}$ and $\uparrow Y = \{x \mid \exists y \in Y.\, y \leq x\}$. In case $Y = \{y\}$ is a singleton we shall write $\downarrow y$ ($\uparrow y$) instead of $\downarrow \{y\}$ ($\uparrow \{y\}$). We also let $\lessdot$ be the relation: $x \lessdot y$ iff $x < y$ and $\forall z \in E.\; x \leq z \leq y$ implies $x = z$ or $z = y$.

A *trace* (over (Σ, I)) is a Σ-labelled poset $T = (E, \leq, \lambda)$ satisfying:

(T1) $\forall e \in E.\quad \downarrow e$ is a finite set
(T2) $\forall e, e' \in E.\quad e \lessdot e' \Rightarrow \lambda(e) D \lambda(e')$.
(T3) $\forall e, e' \in E.\quad \lambda(e) D \lambda(e') \Rightarrow e \leq e'$ or $e' \leq e$.

We shall refer to members of E as *events*. The trace $T = (E, \leq, \lambda)$ is said to be *finite* if E is a finite set. Otherwise it is an infinite trace. Note that E is always a countable set. T is said to be *non-empty* in case $E \neq \emptyset$. We let $TR^{fin}(\Sigma, I)$ be the set of finite traces and $TR^{inf}(\Sigma, I)$ be the set of infinite traces over (Σ, I) and set $TR(\Sigma, I) = TR^{fin}(\Sigma, I) \cup TR^{inf}(\Sigma, I)$. Often we will write TR^{fin} instead of $TR^{fin}(\Sigma, I)$ etc.

Let $T = (E, \leq, \lambda)$ be a trace. A *configuration* is a finite subset $c \subseteq E$ such that $c = \downarrow c$. We let C_T be the set of configurations of T and let c, c', c'' range over C_T. Note that $\emptyset$, the empty set, is a configuration and $\downarrow e$ is a configuration for every $e \in E$. Finally, the transition relation $\to_T \subseteq C_T \times \Sigma \times C_T$ is given by: $c \stackrel{a}{\to}_T c'$ iff there exists $e \in E$ such that $\lambda(e) = a$ and $e \notin c$ and $c' = c \cup \{e\}$. It is easy to see that if $c \stackrel{a}{\to}_T c'$ and $c \stackrel{a}{\to}_T c''$ then $c' = c''$.

3 The Main Result

The first order theory of traces is formulated by assuming a countable set of individual variables $Var = \{x, y, z, \dots\}$; a family of unary predicates $\{R_a\}_{a \in \Sigma}$; a binary predicate $\leq$. Then $FO(\Sigma, I)$, the set of formulas in the first order theory of traces (over (Σ, I)), is given by the syntax:

$$FO(\Sigma, I) ::= R_a(x) \mid x \leq y \mid \sim \varphi \mid \varphi \vee \varphi' \mid (\exists x)\varphi$$

Thus the syntax does not explicitly involve I. However, it is reflected in $\leq$ that will be interpreted as the partial order relation associated with a trace which does indeed respect the independence relation I.

Given a trace $T = (E, \leq, \lambda)$ and an associated valuation $V : Var \rightarrow E$, the relation $T \models_V^{FO} \varphi$ will denote that T is a model of $\varphi \in FO(\Sigma, I)$. This notion is defined in the expected manner. In particular, $T \models_V^{FO} R_a(x)$ iff $\lambda(V(x)) = a$ and $T \models_V^{FO} x \leq y$ iff $V(x) \leq V(y)$. As usual, a sentence is a formula with no free variables. L_φ will denote the set of models of the sentence φ: $L_\varphi = \{T \mid T \in TR \text{ and } T \models^{FO} \varphi\}$

We will say that $L \subseteq TR$ is FO-definable iff there exists a sentence $\varphi \in FO(\Sigma, I)$ such that $L = L_\varphi$.

The set of formulas of our linear time temporal logic of traces (*LTrL*) is defined as follows:

$$LTrL(\Sigma, I) ::= \underline{tt} \mid \sim \alpha \mid \alpha \vee \beta \mid \langle a \rangle \alpha \mid \alpha U \beta \mid \langle a^{-1} \rangle \underline{tt}$$

Thus the next state modality is indexed by actions. There is also a very restricted version of the previous state modality. Indeed the number of past formulas is bounded by the size of Σ. For achieving the present aims, there is no need for atomic propositions. It is worth mentioning that if atomic propositions are to be introduced then the valuations must be required to respect the independence relation in a suitable fashion. The logic will become undecidable otherwise [11]. In the current framework, a model of *LTrL* is just a trace $T = (E, \leq, \lambda)$. The relation $T, c \models \alpha$ will denote that $\alpha \in LTrL(\Sigma, I)$ is satisfied at the configuration $c \in C_T$. This notion is defined via:

- $T, c \models \underline{tt}$. Furthermore $\sim$ and $\vee$ are interpreted in the usual way.
- $T, c \models \langle a \rangle \alpha$ iff $\exists c' \in C_T.\ c \xrightarrow{a}_T c'$ and $T, c' \models \alpha$.
- $T, c \models \alpha U \beta$ iff $\exists c' \in C_T.\ c \subseteq c'$ and $T, c' \models \beta$ and $\forall c'' \in C_T.\ c \subseteq c'' \subset c'$ implies $T, c'' \models \alpha$.
- $T, c \models \langle a^{-1} \rangle \underline{tt}$ iff $\exists c' \in C_T.\ c' \xrightarrow{a}_T c$.

The derived "sometime" and "always" modalities have pleasant semantics. More precisely, with $\Diamond \alpha \overset{\Delta}{\Leftrightarrow} \underline{tt}\, U \alpha$ and $\Box \alpha \overset{\Delta}{\Leftrightarrow} \sim \Diamond \sim \alpha$, we have: $T, c \models \Box \alpha$ iff $\forall c' \in C_T.\ c \subseteq c'$ implies $T, c' \models \alpha$. Thus arbitrary liveness and safety properties interpreted over the global states of a distributed system can be formulated in *LTrL*. With each formula $\alpha \in LTrL(\Sigma, I)$, we can associate a set of traces as follows: $L_\alpha = \{T \in TR \mid T, \emptyset \models \alpha\}$. We say that $L \subseteq TR$ is *LTrL*-definable iff there exists a formula $\alpha \in LTrL(\Sigma, I)$ such that $L = L_\alpha$. Our main result can now be stated.

Theorem 1
Let $L \subseteq TR^{\inf}$. Then L is FO-definable iff L is $LTrL$-definable.

Indeed this result goes through in case $L \subseteq TR^{fin}$ or $L \subseteq TR$. We note that in case $I = \emptyset$, Theorem 1 is just the expressiveness result of [6] in a different and slightly weakened (because of the past modalities) form. As the first order theory of traces is decidable [4] and our translations are constructive we immediately obtain:

Corollary 2 The satisfiability problem for *LTrL* is decidable.

To bring out one more consequence of Theorem 1, we shall define $LTL(\Sigma)$, linear time temporal logic interpreted over Σ-sequences. We will use Σ^* and Σ^ω to denote the set of finite and infinite sequences over Σ respectively. We will use Σ^∞ for $\Sigma^* \cup \Sigma^\omega$.

The syntax of $LTL(\Sigma)$ is given by:

$$LTL(\Sigma) ::= \underline{tt} \mid \sim \widehat{\alpha} \mid \widehat{\alpha} \vee \widehat{\beta} \mid \langle a \rangle \widehat{\alpha} \mid \widehat{\alpha}\ \mathcal{U}\ \widehat{\beta}.$$

For $\sigma \in \Sigma^\infty$, let $prf(\sigma)$ denote the set of finite prefixes of σ and let $\tau \sqsubseteq \tau'$ denote that τ is a prefix of τ'. Then $\sigma, \tau \models \hat{\alpha}$ will stand for $\hat{\alpha}$ being satisfied at the prefix τ of σ. This notion is defined in the usual way.

- $\sigma, \tau \models \underline{tt}$. The connectives $\sim$ and $\vee$ are interpreted in the standard fashion.
- $\sigma, \tau \models \langle a \rangle \hat{\alpha}$ iff $\tau a \in prf(\sigma)$ and $\sigma, \tau a \models \hat{\alpha}$.
- $\sigma, \tau \models \hat{\alpha} \mathcal{U} \hat{\beta}$ iff $\exists \tau' \in prf(\sigma)$ such that $\tau \sqsubseteq \tau'$ and $\sigma, \tau' \models \hat{\beta}$. Moreover for every $\tau'' \in prf(\sigma)$, if $\tau \sqsubseteq \tau'' \sqsubset \tau'$ then $\sigma, \tau'' \models \hat{\alpha}$.

Next, let $T = (E, \leq, \lambda) \in TR$. Then $\sigma \in \Sigma^\infty$ is a *linearisation* of T iff there exists a map $\rho : prf(\sigma) \rightarrow C_T$, such that, the following conditions are met:

(i) $\rho(\epsilon) = \emptyset$ (ϵ is the null string)

(ii) $\forall \tau a \in prf(\sigma)$ with $\tau \in \Sigma^*$, $\rho(\tau) \xrightarrow{a}_T \rho(\tau a)$

(iii) $\forall e \in E\ \exists \tau \in prf(\sigma).\ e \in \rho(\tau)$.

The function ρ will be called a *run map* of the linearisation σ. Note that the run map of a linearisation is unique. In what follows we shall let $\mathrm{lin}(T)$ to be the set of linearisations of the trace T. The notion of linearisation induces the well-known equivalence relation $\approx_I \subseteq \Sigma^\infty \times \Sigma^\infty$ via: $\sigma \approx_I \sigma'$ iff there exists a trace T, such that, $\sigma, \sigma' \in \mathrm{lin}(T)$. A formula $\hat{\alpha}$ is said to be *trace consistent* if $\sigma, \varepsilon \models \hat{\alpha}$ and $\sigma \approx_I \sigma'$ implies $\sigma', \varepsilon \models \hat{\alpha}$; for every $\sigma, \sigma' \in \Sigma^\infty$. As mentioned earlier, specifications that are formulated as trace consistent formulas can be often verified efficiently using partial order reduction techniques. *LTrL* provides a characterisation of trace consistent *LTL* formulas in the following sense.

Corollary 3 For every formula $\alpha \in LTrL(\Sigma, I)$ there is a trace consistent formula $\hat{\alpha} \in LTL(\Sigma)$, s.t. $\bigcup\{\mathrm{lin}(T)|T, \emptyset \vDash \alpha\} = \{\sigma|\sigma, \varepsilon \vDash \hat{\alpha}\}$. For every trace consistent $LTL(\Sigma)$ formula $\hat{\alpha}$ there is a $LTrL(\Sigma, I)$ formula α such that $\{\sigma|\sigma, \varepsilon \vDash \hat{\alpha}\} = \bigcup\{\mathrm{lin}(T)|\ T, \emptyset \vDash \alpha\}$.

4 The Proof

The structure of the proof of Theorem 1 can be brought out by breaking it up into the following steps.

Lemma 4 Let $\alpha \in LTrL(\Sigma, I)$. Then there exists $\varphi \in FO(\Sigma, I)$ such that for every $T \in TR^{inf}$: $T, \emptyset \models \alpha$ iff $T \models^{FO} \varphi$.

Proof
The key observation underlying the proof is that a configuration can be described in $FO(\Sigma, I)$ in terms of its maximal elements. There can be no more than $|\Sigma|$ maximal elements in a configuration.

In $FO(\Sigma, I)$ the variables range over events, but we can use a finite set of variables to represent a configuration. Intuitively a set of variables X represents in a given valuation $V : Var \to E$ the configuration $c_V^X = \{e \mid \exists z \in X.\ e \leq V(z)\}$.

For every set of variables X and every formula α of *LTrL* we will construct a formula φ_α^X of $FO(\Sigma, I)$ with free variables in the set X. This formula will have the property that for every valuation $V : Var \to E$:

$$T \vDash_V^{FO} \varphi_\alpha^X \quad \text{iff} \quad T, c_V^X \vDash \alpha \qquad (1)$$

In particular taking $X = \emptyset$ we will obtain the thesis of the lemma.

The construction proceeds by induction on α. If $\alpha = \underline{\mathrm{tt}}$ then for every X we put $\varphi_\alpha^X = \forall\ z.\ (z \leq z)$. The cases for disjunction and negation are also obvious.

Suppose $\alpha = \langle a \rangle \beta$. Let $X = \{x_1, \ldots, x_k\}$ (this set may be empty). We let φ_α^X to be:

$$\exists y.\ R_a(y) \wedge \varphi_\beta^{X \cup \{y\}} \wedge \Big(\bigwedge_{i=1,\ldots,k} y \not\leq x_i \Big) \wedge \Big(\forall z.\ z < y \Rightarrow \bigvee_{i=1,\ldots,k} z \leq x_i \Big)$$

Suppose $\alpha = \beta U \gamma$. First, for two sets of variables Y, Z we define the formulas

$$\mathrm{Below}(Y, Z) = \bigwedge_{y \in Y} \bigvee_{z \in Z} y \leq z$$
$$\mathrm{SBelow}(Y, Z) = \mathrm{Below}(Y, Z) \wedge \neg\,\mathrm{Below}(Z, Y)$$

Intuitively formula $\mathrm{Below}(X, Y)$ says that all the events in the configuration represented by Y belong to a configuration represented by Z. The formula $\mathrm{SBelow}(X, Y)$ says the same plus the fact that the configurations are not equal. With the help of this formula we define φ_α^X for $X \neq \emptyset$ by:

$$\exists Z.\ \mathrm{Below}(X, Z) \wedge \varphi_\gamma^Z \wedge \forall Y.\ (\mathrm{Below}(X, Y) \wedge \mathrm{SBelow}(Y, Z)) \Rightarrow \varphi_\beta^Y$$

The quantifier $\exists Z$ is an abbreviation of $\exists z_1, \ldots, \exists z_{|\Sigma|}$. Similarly for $\forall Y$. We let $\varphi_\alpha^\emptyset$ to be:

$$\varphi_\gamma^\emptyset \vee \exists Z.\ \varphi_\gamma^Z \wedge \varphi_\beta^\emptyset \wedge \forall Y.\ \mathrm{SBelow}(Y, Z) \Rightarrow \varphi_\beta^Y$$

Finally, if $\alpha = \langle a^{-1} \rangle \underline{\mathrm{tt}}$ then the formula φ_α^X is

$$\bigvee_{x \in X} \Big(R_a(x) \ \wedge \ \bigwedge_{x' \in X} x \neq x' \Rightarrow x \not\leq x' \Big)$$

By induction on α one can show that the condition (1) is satisfied. □

The other direction is much more difficult. Our goal is:

Lemma 5 Let $\varphi \in FO(\Sigma, I)$. Then there exists $\alpha \in LTrL(\Sigma, I)$ such that for every $T \in TR^{\mathrm{inf}}$: $T \models^{FO} \varphi$ iff $T, \emptyset \models \alpha$.

The line of the proof is as follows. First we will define a decomposition of traces into traces with special properties. Next we will show the above lemma for each of these special traces. Finally we will put the obtained formulas together using properties of our decomposition.

4.1 Decomposition of traces

Our decomposition is done in two steps. First a trace is split into finite and infinite part. Then the infinite part turns out to be a disjoint union of infinite traces and we separate the components of this part.

Let $T = (E, \leq, \lambda)$ be a trace. Then $alph(T) = \{\lambda(e) \mid e \in E\}$. Denote $\Sigma_T^{fin} = \{a \mid \lambda^{-1}(a) \text{ is a finite set}\}$. The trace T is called *perpetual* if it is non-empty and $\Sigma_T^{fin} = \emptyset$. Hence every perpetual trace is infinite but converse is not always true. The trace T is called *directed* iff every two events $e_1, e_2 \in E$ have an upper bound under $\leq$, i.e., there exists e, such that, $e_1 \leq e$ and $e_2 \leq e$.

We now define the Σ-labelled posets $\mathrm{fin}(T)$ and $\inf(T)$ via:

$$\mathrm{fin}(T) = (E_{fin}, \leq_{fin}, \lambda_{fin}) \text{ and } \inf(T) = (E_{inf}, \leq_{inf}, \lambda_{inf})$$

where $E_{fin} = \{e \mid \exists e'.\ e \leq e' \text{ and } \lambda(e') \in \Sigma_T^{fin}\}$ and $E_{inf} = E - E_{fin}$. Furthermore, $\leq_{fin}$ ($\leq_{inf}$) is $\leq$ restricted to $E_{fin} \times E_{fin}$ ($E_{inf} \times E_{inf}$) and λ_{fin} (λ_{inf}) is λ restricted to E_{fin} (E_{inf}). The following observation follows easily from the definitions.

Proposition 6 For every trace T, $\text{fin}(T)$ is a finite trace. Further, $\inf(T)$ is a perpetual trace iff T is an infinite trace.

Next we decompose $\inf(T)$.

Proposition 7 Let $T = (E, \leq, \lambda)$ be a perpetual trace. Then there exists a unique family of traces $\{T_i = (E_i, \leq_i, \lambda_i)\}_{i=1}^m$ with $m \leq |\Sigma|$ such that the following conditions are satisfied:

(i) Each T_i is a perpetual directed trace.

(ii) For each $i, j \in \{1, \ldots, m\}$, if $i \neq j$ then $E_i \cap E_j = \emptyset$ and $alph(T_i) \times alph(T_j) \subseteq I$.

(iii) $E = \cup_{i=1}^m E_i$, $\leq = \cup_{i=1}^m \leq_i$ and $\lambda = \cup_{i=1}^m \lambda_i$.

Proof
Let $T = (E, \leq \lambda)$ be a perpetual trace and let $D_T = (alph(T) \times alph(T)) \cap D$. Define a binary relation $\leftrightarrow \subseteq E \times E$ via:

$$e \leftrightarrow e' \quad \text{iff} \quad \exists e''.\ e \leq e'' \text{ and } e' \leq e''. \tag{2}$$

We wish to show that $\leftrightarrow$ is an equivalence relation. For this we will need three observations.

Observation 7.1 Suppose $(a, b) \in D_T$ and $e \in E$ with $\lambda(e) = a$. Then there exists $e' \geq e$ with $\lambda(e') = b$.

To see this, note that as T is perpetual, there must exist infinitely many events labelled by b. For each such event e_b we have $e_b \leq e$ or $e \leq e_b$ by condition T3 in the definition of a trace.

It cannot be the case that all these events are $\leq$-smaller than e; this would contradict the condition (T1) of the definition of a trace. Hence there is an event e' labelled by b that is not $\leq$-smaller than e. By the condition (T3) we have: $e \leq e'$.

Observation 7.2 Let $e, e' \in E$ with $e < e'$. Then $(\lambda(e), \lambda(e')) \in D_T^*$.

As might be expected, D_T^* is the (reflexive and) transitive closure of the relation D_T. Let us prove Observation 7.2. Call a *path* from e to e' in T a sequence $e = e_0 \lessdot e_1 \lessdot \cdots \lessdot e_n = e'$. Clearly such a path must exist because $e < e'$. This follows from $T1$ in the definition of a trace. Again, by condition $T2$ in the definition of a trace, we have $(\lambda(e_i), \lambda(e_{i+1})) \in D_T$ for $0 \leq i < n$.

Observation 7.3 For every $e, e' \in E$ we have: $e \leftrightarrow e'$ iff $(\lambda(e), \lambda(e')) \in D_T^*$

If this observation holds then $\leftrightarrow$ is an equivalence relation because D_T^* is an equivalence relation. To establish the observation first assume that $e'' \in E$ with $e \leq e''$ and $e' \leq e''$ so that $e \leftrightarrow e'$. From Observation 7.2 and the fact that D_T^* is an equivalence relation, we at once have $(\lambda(e), \lambda(e')) \in D_T^*$. So next assume that $(\lambda(e), \lambda(e')) \in D_T^*$ with $\lambda(e) = a$ and $\lambda(e') = b$. If $a = b$ then $e \leftrightarrow e'$ follows at once from condition $T3$ in the definition of a trace. So assume $a \neq b$. Let $a_0, a_1, \ldots, a_n$ be a sequence such that $a = a_0$, $a_n = b$ and $(a_i, a_{i+1}) \in D_T$ for $0 \leq i < n$. By repeated applications of Observation 7.1 we can find a sequence of events $e_0, e_1, \ldots, e_n$ in E such that $e = e_0$, $\lambda(e_i) = a_i$ and $e_i \leq e_{i+1}$ for $0 \leq i < n$. Since $\lambda(e_n) = b = \lambda(e')$ we must have $e' \leq e_n$ or $e_n \leq e'$. In either case, $e \leftrightarrow e'$ as required.

To finish the proof of the proposition, let $\{eq_1, eq_2, \ldots, eq_m\}$ be the set of D_T^*- equivalence classes of $alph(T)$. Define $T_i = (T|eq_i, \leq |eq_i, \lambda|eq_i)$ where $|eq_i$ denotes the restriction to the events labelled with the letters in eq_i. Conditions (i) and (ii) follow from Observation 7.3. Condition (iii) follows directly from the definition of the traces T_i. $\square$

Definition 8 (Shape) The *shape* of a perpetual trace T is the family $\{alph(T_i)\}_{i=1}^m$ where $\{T_i\}_{i=1}^m$ is the decomposition described above. (In other words the shape of T is the D_T^*- equivalence classes of $alph(T)$)

A family $\{\Sigma_i\}_{i=1}^m$ is a *shape in an alphabet* (Σ, I) if it is the shape of some perpetual trace over this alphabet.

4.2 Decomposing formulas in FO

Here we show a couple of composition lemmas which will allow us to reason about the properties of the whole trace in terms of the properties of its components. Before doing this, let us recall, for the sake of completeness, an easy case of composition theorem of Feferman and Vaught [5]. The reader familiar with this topic can proceed directly to Lemma 10.

Let us fix some finite relational signature $\text{Sig} = \{R_1, \ldots, R_l\}$. Given two structures

$$\mathcal{A} = \langle A, R_1^{\mathcal{A}}, \ldots, R_l^{\mathcal{A}} \rangle \qquad \mathcal{B} = \langle B, R_1^{\mathcal{B}}, \ldots, R_l^{\mathcal{B}} \rangle$$

of this signature we define their *disjoint union* as the structure $\mathcal{A} \oplus \mathcal{B}$ of the signature $\text{Sig} \cup \{in_1, in_2\}$:

$$\mathcal{A} \oplus \mathcal{B} = \langle A \oplus B, R_1^{\mathcal{A}} \oplus R_1^{\mathcal{B}}, \ldots, R_l^{\mathcal{A}} \oplus R_l^{\mathcal{B}}, in_1^{\mathcal{A}\oplus\mathcal{B}}, in_2^{\mathcal{A}\oplus\mathcal{B}} \rangle$$

here $A \oplus B$ and $R_i^{\mathcal{A}} \oplus R_i^{\mathcal{B}}$ stand for disjoint sums of the appropriate sets and $in_1^{\mathcal{A}\oplus\mathcal{B}}(a)$ holds if $a \in A$. Similarly $in_2^{\mathcal{A}\oplus\mathcal{B}}(b)$ holds if $b \in B$.

Theorem 9 (Composition thm. for disjoint sum) *Let* Sig *be a finite relational signature. Let φ be a sentence of $FO(\mathrm{Sig} \cup \{in_1, in_2\})$. There exists a finite collection of pairs $(\psi_1, \psi'_1), \dots, (\psi_k, \psi'_k)$ of $FO(\mathrm{Sig})$ sentences, such that, for every two structures $\mathcal{A}$, $\mathcal{B}$ of the signature* Sig *we have:*

$$\mathcal{A} \oplus \mathcal{B} \vDash \varphi \text{ iff there exists } i \in \{1, 2, \dots, k\} \text{ with } \mathcal{A} \vDash \psi_i \text{ and } \mathcal{B} \vDash \psi'_i.$$

Proof

The proof is a standard application of Ehrenfeucht-Fraïssé games. For description of the games see for example [2]. We denote the n-move game on structures $\mathcal{A}$ and $\mathcal{B}$ by $G_n(\mathcal{A}, \mathcal{B})$. Let us denote by $\mathsf{qd}(\theta)$ the *quantifier depth* of the sentence θ. We define an n-theory of a structure $\mathcal{C}$ as the set of sentences $\mathrm{Th}_n(\mathcal{C}) = \{\theta : \mathsf{qd}(\theta) \leq n \text{ and } \mathcal{C} \vDash \theta\}$. We have the following characterisation of n-theories in terms of Ehrenfeucht-Fraïssé games

Observation 9.1 Two structures $\mathcal{A}$, $\mathcal{B}$ have the same n-theories iff Duplicator has a winning strategy in the n-move Ehrenfeucht-Fraïssé game. Every n-theory is equivalent to a single sentence, i.e., for every n-theory Γ there exist a sentence θ_Γ such that for every structure $\mathcal{A}$: $\mathrm{Th}_n(\mathcal{A}) = \Gamma$ iff $\mathcal{A} \vDash \theta_\Gamma$.

The proof of this observation relies on the fact that the signatures are finite and relational.

The next observation is that the n-theory of $\mathcal{A} \oplus \mathcal{B}$ is determined by the n-theories of $\mathcal{A}$ and $\mathcal{B}$. Indeed suppose that $\mathrm{Th}_n(\mathcal{A}) = \mathrm{Th}_n(\mathcal{A}')$ and $\mathrm{Th}_n(\mathcal{B}) = \mathrm{Th}_n(\mathcal{B}')$. By Observation 9.1 it is enough to show that Duplicator has a winning strategy in the n-move game $G_n(\mathcal{A} \oplus \mathcal{B}, \mathcal{A}' \oplus \mathcal{B}')$. By assumption Duplicator has winning strategies in the games $G_n(\mathcal{A}, \mathcal{A}')$ and $G_n(\mathcal{B}, \mathcal{B}')$. The strategy in $G_n(\mathcal{A} \oplus \mathcal{B}, \mathcal{A}' \oplus \mathcal{B}')$ is to copy moves of Spoiler in this game to $G_n(\mathcal{A}, \mathcal{A}')$ or $G_n(\mathcal{B}, \mathcal{B}')$ and consult the strategies there. For example if Spoiler puts a pebble on some element of the $\mathcal{A}$ component of $\mathcal{A} \oplus \mathcal{B}$ then we put Spoilers pebble on the same element in the game $G_n(\mathcal{A}, \mathcal{A}')$. The winning strategy of Duplicator in this game puts a pebble on some element of $\mathcal{A}'$ and we copy this move by putting a pebble on the same element of the $\mathcal{A}'$ component of $\mathcal{A}' \oplus \mathcal{B}'$. It should be clear that such a strategy is winning for Duplicator.

After these preliminary remarks we are ready to prove the theorem. Let φ be a $FO(\mathrm{Sig} \cup \{in_1, in_2\})$ sentence. Let n be the quantifier depth of φ. Let $(\Gamma_1, \Gamma'_1), \dots, (\Gamma_k, \Gamma'_k)$ be all pairs of n-theories such that:

$$\text{if } \mathrm{Th}_n(\mathcal{A}) = \Gamma_i \text{ and } \mathrm{Th}_n(\mathcal{B}) = \Gamma'_i \text{ then } \varphi \in \mathrm{Th}_n(\mathcal{A} \oplus \mathcal{B})$$

The number of such pairs is finite because it can be proved by simple induction on n that there are finitely many n-theories. From Observation 9.1 we know that for every Γ_i there exists a formula ψ_i, such that, for every structure $\mathcal{A}$: $\mathrm{Th}_n(\mathcal{A}) = \Gamma_i$ iff $\mathcal{A} \vDash \psi_i$. Similarly for every Γ'_i we can find ψ'_i. We claim that $(\psi_1, \psi'_1), \dots (\psi_k, \psi'_k)$ satisfies the thesis of the theorem.

For left to right implication suppose that $\mathcal{A} \oplus \mathcal{B} \vDash \varphi$. Then $\varphi \in \mathrm{Th}_n(\mathcal{A} \oplus \mathcal{B})$. Hence there exists i, s.t. $\mathrm{Th}_n(\mathcal{A}) = \Gamma_i$ and $\mathrm{Th}_n(\mathcal{B}) = \Gamma'_i$. So $\mathcal{A} \vDash \psi_i$ and $\mathcal{B} \vDash \psi'_i$. The proof of the reverse implication is similar. □

Let us now come back to decomposing traces. First we show that we can separate finite and infinite part.

Lemma 10 Let $\varphi \in FO(\Sigma, I)$. Then there exists a finite collection of pairs $(\psi_1, \psi'_1), (\psi_2, \psi'_2), \dots, (\psi_k, \psi'_k)$, such that, $\psi_i, \psi'_i \in FO(\Sigma, I)$, for each i, and for every $T \in TR^{inf}$: $T \models^{FO} \varphi$ iff there is $i \in \{1, 2, \dots, k\}$ with $\mathrm{fin}(T) \models^{FO} \psi_i$ and $\mathrm{inf}(T) \models^{FO} \psi'_i$.

Proof

Let $\varphi \in FO(\Sigma, I)$ be given. We claim that there exists a formula φ', such that, for every infinite trace T:

$$T \vDash \varphi \quad \text{iff} \quad \mathrm{fin}(T) \oplus \mathrm{inf}(T) \vDash \varphi' \tag{3}$$

For this we show that in $\mathrm{fin}(T) \oplus \mathrm{inf}(T)$ we can recover the ordering from T by means of a first order formula. Recall that $\mathrm{fin}(T) \oplus \mathrm{inf}(T)$ is a structure of a signature $\{R_a\}_{a\in\Sigma} \cup \{\leq, in_1, in_2\}$. The carriers of T and $\mathrm{fin}(T) \oplus \mathrm{inf}(T)$ are the same. Also the interpretations of the relations $\{R_a\}_{a\in\Sigma}$ are the same. The interpretation of $\leq$ relation in $\mathrm{fin}(T) \oplus \mathrm{inf}(T)$ is the (disjoint) union of $\leq_{\mathrm{fin}}$ and $\leq_{\mathrm{inf}}$ where $\mathrm{fin}(T) = (E_{\mathrm{fin}}, \leq_{\mathrm{fin}}, \lambda_{\mathrm{fin}})$ and $\mathrm{inf}(T) = (E_{\mathrm{inf}}, \leq_{\mathrm{inf}}, \lambda_{\mathrm{inf}})$. Consider the formula:

$$\begin{aligned}\theta(x,y) = &\Big(in_1(x) \wedge in_1(y) \wedge x \leq y\Big)\\ &\vee \Big(in_2(x) \wedge in_2(y) \wedge x \leq y\Big)\\ &\vee \Big(in_1(x) \wedge in_2(y) \wedge \exists z_1 \exists z_2.\ in_1(z_1)\\ &\qquad \wedge in_2(z_2) \wedge D(z_1, z_2) \wedge x \leq z_1 \wedge z_2 \leq y\Big)\end{aligned}$$

where $D(z_1, z_2)$ is a formula stating that the labels of z_1 and z_2 are dependent. It is not difficult to check that for all nodes x, y of T we have: $T \vDash x \leq y$ iff $\mathrm{fin}(T) \oplus \mathrm{inf}(T) \vDash \theta(x, y)$. Hence taking φ and replacing all subformulas of the form $x \leq y$ by $\theta(x, y)$ we obtain a formula φ' satisfying the condition (3). The thesis of the lemma follows directly from Theorem 9. $\square$

Next we further break up the assertions concerning $\inf(T)$ to mimic the decomposition described in Proposition 7.

Lemma 11 Let $\varphi \in FO(\Sigma, I)$ and $sh = \{\Sigma_i\}_{i=1}^m$ be a shape of (Σ, I). Then there exists a finite array of formulas

$$(\theta_1^1, \ldots, \theta_m^1), (\theta_1^2, \ldots, \theta_m^2), \ldots, (\theta_1^n, \ldots, \theta_m^n)$$

such that the following conditions are satisfied:

(i) $\theta_i^j \in FO(\Sigma_i, I)$ for every $i \in \{1, 2, \ldots, m\}$ and every $j \in \{1, 2, \ldots, n\}$. (Observe that the formulas with different subscripts have disjoint alphabets.)

(ii) Suppose $T \in TR^{inf}$, and $\inf(T)$ is of shape sh. Let $\{T_i\}_{i=1}^m$ be a decomposition of $\inf(T)$ as in Proposition 7. We have that $\inf(T) \models^{FO} \varphi$ iff there exists $j \in \{1, 2, \ldots, n\}$, s.t., $T_i \models^{FO} \theta_i^j$ for all $i = 1, \ldots, m$.

This lemma follows easily from Proposition 7 and another easy application of Theorem 9.

4.3 Translation for components

Here we present a translation of $FO(\Sigma, I)$ formulas that works for finite traces as well as for perpetual directed traces.

Lemma 12 Let $\varphi \in FO(\Sigma, I)$. Then there exists a formula $\alpha \in LTrL(\Sigma, I)$ such that for every finite or perpetual directed $T \in TR$ we have: $T \models^{FO} \varphi$ iff $T, \emptyset \models \alpha$.

This lemma constitutes the heart of the proof of Theorem 1. The main ingredients involved in establishing the lemma are an observation made in [4] and a new normal form linearisation of traces. The first step is:

Lemma 13 Let $\varphi \in FO(\Sigma, I)$. Then there exists a *trace consistent* $\hat{\alpha} \in LTL(\Sigma)$ such that for every $T \in TR$, $T \models^{FO} \varphi$ iff $\sigma, \varepsilon \models \hat{\alpha}$ for *some* linearisation σ of T.

Proof

As observed in a slightly different setting in [4], this lemma follows easily from [6]. To see this, let $FO(\Sigma)$ be the first order theory whose syntax is exactly that of $FO(\Sigma, I)$ but whose structures are elements of Σ^∞ with the usual semantics (see for instance [20]). In what follows, the semantic relation of satisfiability associated with the sentences of $FO(\Sigma)$ will be denoted $\models^{fo}$. A simple but basic observation essentially due to Wolfgang Thomas [21] can be stated as:

Observation 13.1 For every sentence $\varphi \in FO(\Sigma, I)$ there exists a sentence $\widehat{\varphi} \in FO(\Sigma)$ such that for every trace T, $T \models^{FO} \varphi$ iff $u \models^{fo} \widehat{\varphi}$ for every $u \in lin(T)$.

Recall that $lin(T)$ is the set of linearisations of T. Now let $T = (E, \leq, \lambda)$ be a trace, $u \in lin(T)$ and $\rho : \mathrm{prf}(u) \to C_T$ the associated run map. Suppose that $e \in E$ and $\lambda(e) = a$. Then there exists a unique $\tau a \in \mathrm{prf}(u)$ such that $e \notin \rho(\tau)$ and $e \in \rho(\tau a)$. Let us call this τa the occurrence of e in u. It is not difficult to show that $e < e'$ in T with $e, e' \in E$ iff there exists $\tau_0 a_0, \tau_1 a_1, \ldots, \tau_n a_n \in \mathrm{prf}(u)$ such that the following conditions are satisfied.

- $\tau_0 a_0$ is the occurrence of e and $\tau_n a_n$ is the occurrence e' in u.
- $\tau_0 a_0 \sqsubseteq \tau_1 a_1 \sqsubseteq \ldots \sqsubseteq \tau_n a_n$.
- $1 \leq n \leq |\Sigma|$ and $a_i \; D \; a_{i+1}$ for $0 \leq i < n$.

All these conditions can be expressed in $FO(\Sigma)$ and this easily leads to Observation 13.1.

Now, by the expressiveness result of [6, 23], for each sentence $\widehat{\varphi} \in FO(\Sigma)$ there exists $\widehat{\alpha} \in LTL(\Sigma)$ such that:

$$\{u \in \Sigma^\infty \mid u \models^{fo} \widehat{\varphi}\} = \{u \in \Sigma^\infty \mid u, \varepsilon \models \widehat{\alpha}\}.$$

The lemma now follows at once from the definition of trace consistent formulas. $\square$

We now wish to exhibit a normal linearisation of traces that can be described within $LTrL(\Sigma, I)$. As a result, we will be able to translate each $LTL(\Sigma)$-formula $\widehat{\alpha}$ into $LTrL(\Sigma, I)$-formula α with the property that a trace T satisfies α at a normal configuration iff $\widehat{\alpha}$ is satisfied at the corresponding prefix of the normal linearisation of T. Through the rest of the section we fix a strict linear order $\prec \subseteq \Sigma \times \Sigma$. For $\emptyset \neq \Sigma' \subseteq \Sigma$, $min(\Sigma')$ will denote the least element of Σ' under $\prec$.

Let $T = (E, \leq, \lambda) \in TR$ be a trace. Then the relation $co \subseteq E \times E$ is defined as: $e \; co \; e'$ iff $e \not\leq e'$ and $e' \not\leq e$. Further, for $e, e' \in E$ we set $\Sigma_{ee'} = \lambda(\uparrow e - \uparrow e')$. (For $X \subseteq E$, $\lambda(X) = \{\lambda(x) \mid x \in X\}$.)

Definition 14 Let $T = (E, \leq, \lambda)$ be a trace. Then $lex_T \subseteq E \times E$ is defined as: $e \; lex_T \; e'$ iff $e < e'$ or $e \; co \; e'$ and $min(\Sigma_{ee'}) \prec min(\Sigma_{e'e})$.

Suppose $T = (E, \leq, \lambda)$ is a trace and $e, e' \in E$ with $e \; co \; e'$. Then it is easy to show that $\Sigma_{ee'} \cap \Sigma_{e'e} = \emptyset$ and that both $\Sigma_{ee'}$ and $\Sigma_{e'e}$ are nonempty. Hence lex_T is well-defined.

Lemma 15 Let $T = (E, \leq, \lambda) \in TR$ be a trace. Then (E, lex_T) is a strict linear order.

Proof

Let $e, e' \in E$ with $e \neq e'$. It is straightforward to verify that $e \; lex_T e'$ or $e' \; lex_T \; e$ but not both. So what needs to be shown is that lex_T is transitive.

Let $e_1, e_2, e_3 \in E$ with $e_1 \; lex_T \; e_2$ and $e_2 \; lex_T \; e_3$. To show $e_1 \; lex_T \; e_3$, first note that e_1, e_2 and e_3 must be pairwise distinct. For distinct $i, j \in \{1, 2, 3\}$ we fix (if it exists) an event $e_{ij} \in \uparrow e_i - \uparrow e_j$ labelled with the $\prec$-smallest action among those occurring in $\uparrow e_i - \uparrow e_j$. We need to examine several, quite easy, cases.

Suppose $e_1 < e_2$. Then $\uparrow e_2 - \uparrow e_3 \subseteq \uparrow e_1 - \uparrow e_3$ and $\uparrow e_3 - \uparrow e_1 \subseteq \uparrow e_3 - \uparrow e_2$. As $lex_T(e_2, e_3)$ we get $lex_T(e_1, e_3)$.

The case when $e_2 \leq e_3$ is done similarly. If $e_1 \leq e_3$ then $lex_T(e_1, e_3)$ and we are done.

Suppose $e_1 \, \mathrm{co} \, e_2$ and $e_2 \, \mathrm{co} \, e_3$ and $e_1 \not\leq e_3$. We claim that $e_1 \, \mathrm{co} \, e_3$. If it were $e_3 \leq e_1$ then $\uparrow e_1 - \uparrow e_2 \subseteq \uparrow e_3 - \uparrow e_2$ and $\uparrow e_2 - \uparrow e_3 \subseteq \uparrow e_2 - \uparrow e_1$. Hence $\lambda(e_{32}) \preceq \lambda(e_{12})$ and $\lambda(e_{21}) \preceq \lambda(e_{23})$. We also know that $\lambda(e_{12}) \prec \lambda(e_{21})$. This gives us $\lambda(e_{32}) \prec \lambda(e_{23})$, a contradiction.

Hence we are left with the case when e_1, e_2, e_3 are pairwise in co relation. From $lex_T(e_1, e_2)$ and $lex_T(e_2, e_3)$ we get $\lambda(e_{12}) \prec \lambda(e_{21})$ and $\lambda(e_{23}) \prec \lambda(e_{32})$.

First we claim that:

$$\lambda(e_{13}) \preceq \lambda(e_{12}). \tag{4}$$

Suppose $e_{12} \notin \uparrow e_3$. Then $e_{12} \in \uparrow e_1 - \uparrow e_3$ and (4) follows. So assume that $e_{12} \in \uparrow e_3$. Then $e_{12} \in \uparrow e_3 - \uparrow e_2$. Since $e_2 \; lex_T \; e_3$ we have:

$$\lambda(e_{23}) \prec \lambda(e_{32}) \preceq \lambda(e_{12}). \tag{5}$$

Now we must consider two cases. Suppose $e_{23} \in \uparrow e_1$. Then $e_{23} \in \uparrow e_1 - \uparrow e_3$ and hence $\lambda(e_{13}) \preceq \lambda(e_{23})$ which then leads to (4). Suppose on the other hand $e_{23} \notin \uparrow e_1$. Then $e_{23} \in \uparrow e_2 - \uparrow e_1$ which leads to $\lambda(e_{12}) \prec \lambda(e_{21}) \preceq \lambda(e_{23})$. But from (5) above we now have the contradiction: $\lambda(e_{12}) \prec \lambda(e_{12})$. Hence (4) must hold.

To finish the proof there are two cases to consider. Suppose $e_{31} \in \uparrow e_2$. Then $e_{31} \in \uparrow e_2 - \uparrow e_1$ and from $\lambda(e_{12}) \prec \lambda(e_{21}) \preceq \lambda(e_{31})$ and (4) we can deduce $\lambda(e_{13}) \prec \lambda(e_{31})$. So suppose that $e_{31} \notin \uparrow e_2$. Then $e_{31} \in \uparrow e_3 - \uparrow e_2$ and consequently $\lambda(e_{23}) \prec \lambda(e_{32}) \preceq \lambda(e_{31})$. If $e_{23} \in \uparrow e_1$ then $e_{23} \in \uparrow e_1 - \uparrow e_3$ and hence $\lambda(e_{13}) \preceq \lambda(e_{23}) \prec \lambda(e_{31})$ as desired. If on the other hand, $e_{23} \notin \uparrow e_1$ then $e_{23} \in \uparrow e_2 - \uparrow e_1$ and hence $\lambda(e_{12}) \prec \lambda(e_{21}) \preceq \lambda(e_{23})$. This in turn leads to $\lambda(e_{12}) \prec \lambda(e_{31})$. From (4) we can again conclude that $\lambda(e_{13}) \prec \lambda(e_{31})$.

□

We shall introduce the notion of normal configurations that in turn will enable us to define normal linearisations of traces. Let $T = (E, \leq, \lambda) \in TR$ and $c \in C_T$. Then c is a *normal configuration* iff for every $e \in c$ and every $e' \in E$, if $e' \; lex_T \; e$ then $e' \in c$.

Now, let σ be a linearisation of T with ρ as the run map of σ (as defined in Section 3). Then σ is a *normal linearisation* of T iff $\rho(\tau)$ is a normal configuration for every $\tau \in prf(\sigma)$. It is easy to see that there can be at most one normal linearisation of a trace. Some traces do not have normal linearisations. One of the reasons why we focus on directed perpetual traces is:

Lemma 16 If T is finite or a directed perpetual trace then there exists a unique normal linearisation of T.

Proof

Let c be a configuration of a trace $T = (E, \leq, \lambda)$. We say that the event $e \in E$ is *enabled* at c iff $e \notin c$ and $c \cup \{e\}$ is a configuration. (This notion plays an important role in the proof of the next lemma too). It is easy to see that e is enabled at c iff e is a minimal element of $E - c$ under $\leq$. Next we note that if c is a normal configuration of the trace T and e is the least enabled event at c under lex_T (among all the enabled events at c), then $c \cup \{e\}$ is also a normal configuration. From the fact that the empty configuration is always normal, it now follows that if T is a finite trace then it admits a unique normal linearisation.

One can apply the same reasoning in case of directed perpetual traces but it may be not clear that the obtained sequence contains all the events of the trace. To show that it is indeed the case it is enough to show that for every event e the set $\{e' \mid e' \; lex_T \; e\}$ is finite.

Let $\{a_1, \ldots, a_k\} = alph(T)$. Take an event $e_1 \geq e$ labelled with a_1. Such an event exists because T is directed and perpetual. Then inductively for every $i = 2, \ldots, k$ take $e_i \geq e_{i-1}$ labelled by a_i. We claim that if $e' \; lex_T \; e$ then $e' \leq e_k$. If $e' \leq e$ then it is obvious. If $e' \, \mathrm{co} \, e$ then let a_j be the label of e'. Clearly $e' \leq e_j$. Hence $e' \leq e_k$. We are done, as the set $\downarrow e_k$ is finite. □

The crucial feature of lex_T is that normal configurations are definable in *LTrL*.

Lemma 17 There exists an $LTrL(\Sigma, I)$ formula NRC such that for every $T \in TR$ and every $c \in C_T$: $T, c \models NRC$ iff c is a normal configuration.

Proof
We will say that the event e is at the *top* of the configuration c iff $c \cap \uparrow e = \{e\}$. In other words, e is a maximal element of c under $\leq$. We let $top(c)$ be the set of elements that are on the top of c.

Observation 17.1 Let $T = (E, \leq, \lambda)$ be a trace and $c \in C_T$. Then c is not a normal configuration iff there exist events e, e' and e_1 satisfying the following conditions:

(i) $e \in top(c)$, e' is enabled at c and $e_1 \in \uparrow e' - \uparrow e$,

(ii) $\forall e_2 \in E$, if $e_2 \in \uparrow e - \uparrow e'$ then $\lambda(e_1) \prec \lambda(e_2)$.

To see that this holds assume first that c is not a normal configuration. Then there exists $e_3 \in c$ and $e_3' \notin c$ such that $e_3'\ lex_T\ e_3$. Let $e \in top(c)$ such that $e_3 \leq e$ and let e' be enabled at c such that $e' \leq e_3'$. By the transitivity of lex_T we now have $e'\ lex_T\ e$. Let $x = \min(\Sigma_{e'e})$ and $e_1 \in \uparrow e' - \uparrow e$ such that $\lambda(e_1) = x$. Now suppose $e_2 \in \uparrow e - \uparrow e'$. By the definition of lex_T, we have $x \prec \lambda(e_2)$.

Next suppose there exist e, e' and e_1 fulfilling the conditions specified by Observation 17.1. Let $e_2 \in \uparrow e - \uparrow e'$. Then $\lambda(e_1) \prec \lambda(e_2)$. Hence $min(\Sigma_{e'e}) \preceq \lambda(e_1) \prec \min(\Sigma_{ee'})$. Thus $e'\ lex_T\ e$ and c is not normal.

We need to define an intermediate formula before getting to NRC. In what follows, for $a, b, d \in \Sigma$ let $\Gamma^d_{ab} = \{S \mid S \subseteq \Sigma \text{ and } a \in S \text{ but } b, d \notin S\}$. We will use this notion only in contexts where $a \neq b$ and $a \neq d$.

For $a, b, d \in \Sigma$, define the formula μ^d_{ab} to be $\sim \underline{tt}$ in case $a = b$ or $a = d$. Otherwise,

$$\mu^d_{ab} = \langle b^{-1} \rangle \underline{tt} \wedge \bigvee_{S \in \Gamma^d_{ab}} \alpha_S$$

where

$$\alpha_S = \Big(\bigwedge_{x \in S} \langle x^{-1} \rangle \underline{tt} \Big)\ U\ \Big(\langle d^{-1} \rangle \underline{tt} \wedge \bigwedge_{x \in S} \langle x^{-1} \rangle \underline{tt} \wedge \bigwedge_{y \in (\Sigma - S) - \{d\}} \sim \langle y^{-1} \rangle \underline{tt} \Big).$$

Observation 17.2 Let $T = (E, \leq, \lambda)$ be a trace and $c \in C_T$. Then $T, c \models \mu^d_{ab}$ iff there exist events e, e', e_1 such that the following conditions are satisfied:

(i) $\lambda(e) = a$, $\lambda(e') = b$ and $\lambda(e_1) = d$.

(ii) $e, e' \in top(c)$ with $e \neq e'$ and $e_1 \in \uparrow e' - \uparrow e$.

To see that this must hold first suppose that $T, c \models \mu^d_{ab}$. Then $a \neq b$ and $a \neq d$. Let $S \in \Gamma^d_{ab}$ such that $T, c \models \langle b^{-1} \rangle \underline{tt} \wedge \alpha_S$. Then there exists c' such that $c \subseteq c'$ and

$$T, c' \models \langle d^{-1} \rangle \underline{tt}\ \wedge \bigwedge_{x \in S} \langle x^{-1} \rangle \underline{tt}\ \wedge \bigwedge_{y \in (\Sigma - S) - \{d\}} \sim < y^{-1} > \underline{tt}. \tag{1}$$

Furthermore, for every configuration c'' such that $c \subseteq c'' \subset c'$, we have

$$T, c'' \models \bigwedge_{x \in S} \langle x^{-1} \rangle \underline{tt}. \tag{2}$$

Let $e, e' \in top(c)$ such that $\lambda(e) = a$ and $\lambda(e') = b$. Clearly, $a \neq b$ implies $e \neq e'$.

Now suppose $b = d$. Then by setting $e_1 = e'$, we at once get the desired conclusion. This follows from the fact that $a \neq b$ because $b \notin S$ and hence $\lambda(e) \neq \lambda(e')$. But then $e, e' \in top(c)$ and thus $e' \in \uparrow e' - \uparrow e$.

So assume that $b \neq d$. Then $T, c' \models \sim \langle b^{-1} \rangle \underline{tt}$ because $b \in (\Sigma - S) - \{d\}$.

Let $e_1 \in top(c')$ such that $\lambda(e_1) = d$. Such an e_1 must exist because $T, c' \models \langle d^{-1} \rangle \underline{tt}$. We now wish to argue that e, e' and e_1 have the desired properties.

Let $S = \{a^1, a^2, \ldots, a^k\}$. Note that $a \in S$. Since $T, c \models \bigwedge_{x \in S} \langle x^{-1} \rangle \underline{tt}$, we can fix $e^1, e^2, \ldots e^k \in top(c)$ such that $\lambda(e^j) = a^j$ for each $j \in \{1, 2, \ldots, k\}$. Clearly $e \in \{e^1, \ldots, e^j\}$. We will first argue that $e^j\ co\ e_1$ for every j which will lead to $e\ co\ e_1$. So fix $j \in \{1, 2, \ldots, k\}$ and suppose that $e^j\ co\ e_1$ does not hold. Since $e_1 \in top(c')$ and $c \subseteq c'$ and $e^j \in top(c)$ we can rule out $e_1 < e^j$. So it must be the case that $e^j < e_1$. But this implies that there exists a finite chain $e^j = z_0 \lessdot z_1 \lessdot \cdots \lessdot z_n = e_1$. Since $\lambda(e_1) = d \notin S$ we have $\lambda(e^j) \neq d$ and $n \geq 1$. Let i be the least integer in $\{0, 1, \ldots, n-1\}$ such that $\lambda(z_i) = a^j$ and $\lambda(z_{i+1}) \neq a^j$. Let $\lambda(z_{i+1}) = \widehat{a}$. Clearly $a^j\ D\ \widehat{a}$. Now consider the configuration $\widehat{c} = c\ \cup\ \downarrow z_{i+1}$. It is easy to check that $c \subseteq \widehat{c} \subseteq c'$. Hence $T, \widehat{c} \models \bigwedge_{x \in S} \langle x^{-1} \rangle \underline{tt}$ which then implies $T, \widehat{c} \models \langle (a^j)^{-1} \rangle \underline{tt}$. But $z_{i+1} \in top(\widehat{c})$ and hence $T, \widehat{c} \models \langle (\widehat{a})^{-1} \rangle \underline{tt}$. This is a contradiction because two distinct labels at the top of a configuration can not be in the dependence relation. Thus $e^j\ co\ e_1$ and consequently $e\ co\ e_1$.

Next we must show that $e' \leq e_1$. Since $T, c \models \langle b^{-1} \rangle \underline{tt}$ and $T, c' \models \sim \langle b^{-1} \rangle \underline{tt}$ (recall that we are considering the case $b \neq d$) we know that $e' \notin top(c')$. Hence there exists $e'' \in top(c')$ such that $e' < e''$. If $e'' = e_1$, we are done. Otherwise $e' < e''$ for some $e'' \in top(c')$ with $\lambda(e'') = a^j$ for some $a^j \in S$. We will now argue that this is impossible.

Suppose $e' < e''$ and $\lambda(e'') = a^j \in S$ with $e'' \in top(c')$. Then from $e^j \in top(c)$ and $b \notin S$, we get $a^j\ I\ b$. Consequently $a^j \neq b$. Clearly there exists a non-null path $e' = z_0 \lessdot z_1 \lessdot \cdots \lessdot z_n = e''$. Let i be the largest integer in $\{1,2,\ldots,n\}$ such that $z_i = a^j$ and $z_{i-1} \neq a^j$. Let $\lambda(z_{i-1}) = \widehat{a}$ and $\widehat{c} = c \cup \downarrow z_{i-1}$. It is easy to check that $c \subseteq \widehat{c} \subset c'$ and hence $T, \widehat{c} \models \langle (a^j)^{-1} \rangle \underline{tt}$. But $z_{j-1} \in top(\widehat{c})$ and hence $T, \widehat{c} \models \langle (\widehat{a})^{-1} \rangle \underline{tt}$. We now have a contradiction because $\widehat{a} \neq a^j$ and $\widehat{a}\ D\ a^j$.

To prove the right to left implication of Observation 17.2 assume that the event e, e' and e_1 exist which fulfil the properties specified in the observation. Let $c' = c \cup \downarrow e_1$. Then $e_1 \in top(c')$ and hence $T, c' \models \langle d^{-1} \rangle \underline{tt}$. Let $S = \{\lambda(e'') \mid e'' \in top(c')$ and $e'' \neq e_1\}$. First we assert $a \in S$. This is because $e \in top(c)$ and $e\ co\ e_1$. Hence $e \in top(c')$ as well because $c' = c \cup \downarrow e_1$. By the definition of S we are assured that $d \notin S$. Hence $a \neq d$. Next suppose $b \in S$. Then there exists $e'' \in top(c')$ such that $e'' \neq e_1$ and $\lambda(e'') = b$. But then $e'' \in c \cup \downarrow e_1$ and since $e'' \neq e_1$ implies $e''\ co\ e_1$, we must have $e'' \in c$. In fact $e'' \in top(c)$ because $e'' \in top(c')$ and $c \subseteq c'$. But this implies $e'' = e'$ which contradicts $e' \leq e_1$. Thus $b \notin S$ and consequently $b \neq a$.

Clearly, by the choice of S, we have

$$T, c' \models \langle d^{-1} \rangle \underline{tt} \wedge \bigwedge_{x \in S} \langle x^{-1} \rangle \underline{tt} \wedge \bigwedge_{y \in (\Sigma - S) - \{d\}} \sim \langle y^{-1} \rangle \underline{tt}.$$

It is also clear that $T, c \models \langle b^{-1} \rangle \underline{tt}$. So suppose $c \subseteq c'' \subset c'$ and $\widehat{a} \in S$. Then there exists $e'' \in top(c')$ such that $\lambda(e'') = \widehat{a}$. But then $c' = c \cup \downarrow e_1$ and $d \notin S$ at once leads to $\widehat{a} \in top(c'')$ as well. Hence $T, c'' \models \langle (\widehat{a}^{-1} \rangle \underline{tt}$. We now have $T, c \models \mu^d_{ab}$.

Now we define the desired formula NRC as:

$$NRC = \sim \bigvee_{(a,b) \in I} \langle b \rangle \left(\bigvee_{d \in \Sigma} \left(\mu^d_{ab} \wedge \bigwedge_{d' \prec d} \sim \mu^{d'}_{ba} \right) \right).$$

To see that NRC has the required property assume first that $T = (E, \leq, \lambda)$ is a trace and $\widehat{c} \in C_T$ is a configuration that is not normal. Then by Observation 17.1, there exist event e, e' and e_1 such that $e \in top(\widehat{c})$, e' is enabled at $\widehat{c}$ and $e_1 \in \uparrow e' - \uparrow e$. Further, if $e_2 \in \uparrow e - \uparrow e'$ then $\lambda(e_1) \prec \lambda(e_2)$. Let $\lambda(e) = a$, $\lambda(e') = b$ and $\lambda(e_1) = d$. If $e \leq e'$ then this would lead to $e \leq e_1$ contradicting $e\ co\ e_1$. Hence $e\ co\ e'$ as well. Consequently $a \neq b$ and $a \neq d$. Now consider the configuration $c = \widehat{c} \cup \{e'\}$. Clearly c fulfils the requirements of Observation 17.2 and hence $T, c \models \mu^d_{ab}$. Now suppose $T, c \models \mu^{d'}_{ba}$ for some $d' \prec d$. Then by the definition of the formula $\mu^{d'}_{ba}$ we are assured that $b \neq d'$. Further, we already have $b \neq a$. Now again by Observation 17.2, there exists e_2 such that $e_2 \in \uparrow e - \uparrow e'$ with $\lambda(e_2) \prec \lambda(e_1)$. But this contradicts the criteria justifying the choice of e, e' and e_1. Hence $T, \widehat{c} \models \sim NRC$.

Next suppose $T, \widehat{c} \models \sim NRC$. Then there exists $(a,b) \in I$ and $d \in \Sigma$ such that $T, \widehat{c} \models \langle b \rangle \left(\mu^d_{ab} \wedge \bigwedge_{d' \prec d} \sim \mu^{d'}_{ba} \right)$. Clearly $a \neq b$ and $a \neq d$. Hence there exists an event $e \in top(\widehat{c})$ and an event e' which is enabled at $\widehat{c}$ such that $\lambda(e) = a$ and $\lambda(e') = b$. Moreover with $c = \widehat{c} \cup \{e'\}$, we have $T, c \models \mu^d_{ab} \wedge \bigwedge_{d' \prec d} \sim \mu^{d'}_{ba}$. Because $T, c \models \mu^d_{ab}$ there exists an event e_1 such that $\lambda(e_1) = d$ and $e_1 \in \uparrow e' - \uparrow e$. This follows from Observation 17.2. Now suppose there exists $e_2 \in \uparrow e - \uparrow e'$ such that $\lambda(e_2) = d' \prec d$. If $d' \neq b$ then, by Observation 17.2, we have $T, c' \models \mu^{d'}_{ba}$; a contradiction. Hence it must be the case that $d' = b$ so that $\mu^{d'}_{ba} = \sim \underline{tt}$. But this is again a contradiction, because $\lambda(e_2) = d' = b$ implies that $e' \leq e_2$ or $e_2 \leq e'$ whereas we are supposed to have $e_2\ co\ e'$. Thus $\min(\Sigma_{e'e}) \preceq d \prec \min(\Sigma_{ee'})$. This leads to $e'\ lex_T\ e$, which then guarantees that $\widehat{c}$ is not a normal configuration. $\square$

Using NRC, we now define the map $\|.\| : LTL(\Sigma) \to LTrL(\Sigma, I)$ via:

$$\|\underline{tt}\| = \mathrm{tt} \quad \| \sim \hat{\alpha}\| = \sim \|\hat{\alpha}\| \quad \|\hat{\alpha} \vee \hat{\beta}\| = \|\hat{\alpha}\| \vee \|\hat{\beta}\|$$
$$\|\langle a \rangle \hat{\alpha}\| = \langle a \rangle (NRC \wedge \|\hat{\alpha}\|)$$
$$\|\hat{\alpha}\ \mathcal{U}\ \hat{\beta}\| = (NRC \supset \|\hat{\alpha}\|)\ U\ (NRC \wedge \|\hat{\beta}\|)$$

Lemma 18 Let $\hat{\alpha}$ be a formula of $LTL(\Sigma)$. For every finite or perpetual directed trace T and its normal linearisation σ_0 we have: $\sigma_0, \varepsilon \models \hat{\alpha}$ iff $T, \emptyset \models \|\hat{\alpha}\|$.

The lemma can be obtained without much work by structural induction of $\hat{\alpha}$ using the property of NRC.

We claim that Lemma 12 follows from Lemma 18 and the previous results. To see this, let $\varphi \in FO(\Sigma, I)$ and T be a finite or perpetual directed trace. Then by Lemma 13, there exists $\hat{\alpha} \in LTL(\Sigma)$, such that, $T \models \varphi$ iff $\sigma_0, \varepsilon \models \hat{\alpha}$; where σ_0 is the normal linearisation of T. By Lemma 18: $T \models \varphi$ iff $T, \emptyset \models \|\hat{\alpha}\|$. Thus with each $\varphi \in FO(\Sigma, I)$ we can associate the formula $\|\hat{\alpha}\|$.

4.4 Composing formulas in *LTrL*

Lemma 19 Let $\varphi \in FO(\Sigma, I)$. Then there exists a formula $\alpha \in LTrL(\Sigma, I)$ such that for every $T \in TR^{inf}$, $\inf(T) \models^{FO} \varphi$ iff $\inf(T), \emptyset \models \alpha$.

Proof
Let us fix a $FO(\Sigma, I)$ formula φ. First, for every shape $sh = \{\Sigma_i\}_{i=1}^m$ of (Σ, I) we will construct a *LTrL* formula α_{sh} with the property:

for every trace T with $\inf(T)$ of shape sh:

$$\inf(T) \models^{FO} \varphi \quad \text{iff} \quad \inf(T), \emptyset \models \alpha_{sh} \qquad (6)$$

Let us fix a shape $sh = \{\Sigma_i\}_{i=1}^m$. By Lemma 11, for the shape sh we have an array of $FO(\Sigma, I)$ formulas:

$$(\theta_1^1, \dots, \theta_m^1), (\theta_1^2, \dots, \theta_m^2), \dots, (\theta_1^n, \dots, \theta_m^n) \qquad (7)$$

such that for every trace T, if $\inf(T)$ is of the shape sh and $\{T_i\}_{i=1}^m$ is the decomposition of $\inf(T)$ as in Proposition 7 then:

$$\begin{aligned}\inf(T) \vDash^{FO} \varphi \text{ iff } & \text{there is } j \in \{1, 2, \dots n\} \text{ such that} \\ & \text{for all } i \in \{1, 2, \dots, m\},\ T_i, \emptyset \vDash \theta_i^j\end{aligned}$$

Moreover each θ_i^j is over the alphabet Σ_i.

By Lemma 12, for every θ_i^j we can find a *LTrL* formula α_i^j such that for every perpetual directed trace T' over the alphabet Σ_i we have: $T' \models^{FO} \theta_i^j$ iff $T', \emptyset \vDash \alpha_i^j$. Hence, for a decomposition $\{T_i = (E_i, \leq_i, \lambda_i)\}_{i=1}^m$ as above and for every $j = 1, \dots, n$ we have: $T_i \vDash^{FO} \theta_i^j$ iff $T_i, \emptyset \vDash \alpha_i^j$. Now $\Sigma_i \times \Sigma_j \subseteq I$ whenever $i \neq j$ and $i, j \in \{1, 2, \dots, m\}$. Hence, we are assured that $c \subseteq E$ is a configuration of $\inf(T)$ iff $c_i = c \cap E_i$ is a configuration of T_i for each i. It is easy to establish by structural induction that for every formula γ_i over Σ_i (i.e. γ_i mentioning at most the letters in Σ_i) and for every configuration c of $\inf(T)$, we have $\inf(T), c \models \gamma_i$ iff $T_i, c_i \models \gamma_i$. Since each α_i^j is over the alphabet Σ_i we have: $\inf(T), \emptyset \models \alpha_1^j \wedge \alpha_2^j \cdots \wedge \alpha_m^j$ iff $T_i \models \alpha_i^j$ for each i.

Let us denote $\alpha_1^j \wedge \cdots \wedge \alpha_n^j$ by β^j and let

$$\alpha_{sh} = \beta^1 \vee \cdots \vee \beta^n$$

It should be clear that α_{sh} satisfies the property (6). Next we observe that we can write a formula ν_{sh} in $LTrL$ such that $\inf(T), \emptyset \models \nu_{sh}$ iff $\inf(T)$ is of shape sh. Let $sh = \{\Sigma_i\}_{i=1}^m$ and

$$\Sigma_{sh} = \bigcup_{i=1}^m \Sigma_i.$$

Then

$$\nu_{sh} = \left(\bigwedge_{a \in \Sigma_{sh}} \Diamond \langle a \rangle \underline{tt} \right) \wedge \left(\bigwedge_{a \notin \Sigma_{sh}} \Box \sim \langle a \rangle \underline{tt} \right).$$

Clearly (Σ, I) admits only finitely many shapes. Let SH be the set of all shapes. Then consider the formula:

$$\alpha = \bigwedge_{sh \in SH} (\nu_{sh} \supset \alpha_{sh}).$$

It is not difficult to show now that α satisfies the property required by the lemma. $\Box$

Lemma 20 Let $\alpha_0, \alpha_1 \in LTrL(\Sigma, I)$. Then there exists a formula $\alpha \in LTrL(\Sigma, I)$ such that for every $T \in TR^{inf}$, $T, \emptyset \models \alpha$ iff $\mathrm{fin}(T), \emptyset \models \alpha_0$ and $\inf(T), \emptyset \models \alpha_1$.

Proof
First we define a *LTrL* formula BORDER that holds precisely in the configuration of T consisting of all the events in $\mathrm{fin}(T)$:

$$\begin{aligned}\text{BORDER} \;=\; & (\bigwedge_{a \in \Sigma} < a^{-1} > \underline{tt} \supset \Box \sim < a > \underline{tt}) \wedge \\ & \bigwedge_{b \in \Sigma} \Box(< b > \underline{tt} \supset < b > \Diamond < b > \underline{tt}).\end{aligned}$$

Next we define FIN to be the formula $\Diamond$BORDER. We have that for every trace T: $T, c \vDash FIN$ iff $c \subseteq \mathrm{fin}(T)$ and $T, c \vDash BORDER$ iff $c = E_{fin}$; where E_{fin} is the set of events of $fin(T)$.

Now, with each formula $\alpha \in LTrL$ we associate the formula $fin(\alpha)$ inductively as follows.

$$\begin{aligned} fin(\underline{tt}) = \underline{tt} \qquad & fin(\sim \alpha) = \sim fin(\alpha) \\ fin((\alpha \vee \beta)) &= fin(\alpha) \vee fin(\beta) \\ fin(\langle a \rangle \alpha) &= \langle a \rangle (FIN \wedge fin(\alpha)) \\ fin(\alpha U \beta) &= fin(\alpha)\, U\, (FIN \wedge\ fin(\beta)) \end{aligned}$$

Also, with each formula β we associate the formula $inf(\beta)$ given by: $inf(\beta) = \Diamond(\text{BORDER} \wedge \beta)$. Now, let $T = (E, \leq, \lambda) \in TR^{\inf}$ and $fin(T) = (E_{fin}, \leq_{fin}, \lambda_{fin})$ and $\inf(T) = (E_{\inf}, \leq_{\inf}, \lambda_{\inf})$. It follows from the definitions that $c \subseteq E_{fin}$ is a configuration of $fin(T)$ iff c is a configuration of T. Hence using the properties of the translation map fin defined above we can establish by structural induction on α that $T, c \models fin(\alpha)$ iff $fin(T), c \models \alpha$ for each configuration c of $fin(T)$.

Next we note that $c \subseteq E_{\inf}$ is a configuration of $\inf(T)$ iff $E_{fin} \cup c$ is a configuration of T. Again, by using the property of the map inf, we can show by structural induction, that

$$T, E_{fin} \cup c \models inf(\beta) \text{ iff } \inf(T), c \models \beta$$

for each configuration c of $\inf(T)$. It now follows at once that for every $T \in TR^{inf}, T, \emptyset \models fin(\alpha_0) \wedge inf(\alpha_1)$ iff $\mathrm{fin}(T), \emptyset \models \alpha_0$ and $\inf(T), \emptyset \models \alpha_1$.

$\Box$

Clearly Lemmas 10, 12, and 20 together yield Lemma 5 and we are done.

Using the intermediate lemmas that have been established to prove the main result, it is an easy exercise to derive Corollary 3.

References

[1] R. Alur, D. Peled, and W. Penczek. Model-checking of causality properties. In *LICS '95*, pages 90–100, 1995.

[2] H.-D. Ebbinghaus and J. Flum. *Finite Model Theory*. Springer-Verlag, 1995.

[3] W. Ebinger. *Charakterisierung von Sprachklassen unendlicher Spuren durch Logiken*. PhD thesis, Institut für Informatik, Universität Stuttgart, 1994.

[4] W. Ebinger and A. Muscholl. Logical definability on infinite traces. In *ICALP '93*, volume 700, pages 335–346, 1993.

[5] S. Feferman and R. Vaught. The first order properties of products of algebraic systems. *Fundamenta Mathematicae*, 47:57–103, 1959.

[6] A. Gabbay, A. Pnueli, S. Shelah, and J. Stavi. On the temporal analysis of fairness. In *7th Ann. ACM Symp. on Principles of Programming Languages*, pages 163–173, 1980.

[7] D. Gabbay, I. Hodkinson, and M. Reynolds. *Temporal Logic : Mathematical Foundations and Computational Aspects*, volume 1. Clarendon Press, Oxford, G.B., 1994.

[8] P. Gastin and A. Petit. Asynchronous cellurar automata for infinite traces. In *ICALP '92*, volume 623 of *LNCS*, pages 583–594, 1992.

[9] P. Godefroid. *Partial-order methods for the verification of concurrent systems*, volume 1032 of *LNCS*. Springer-Verlag, 1996.

[10] H. Kamp. *Tense Logic and the Theory of Linear Order*. PhD thesis, University of California, 1968.

[11] M. Mukund and P.S. Thiagarajan. Linear time temporal logics over traces. In *MFCS'96*, volume 1113 of *LNCS*, pages 62–92, 1996.

[12] P. Niebert. A ν-calculus with local views for sequential agents. In *MFCS '95*, volume 969 of *LNCS*, pages 563–573, 1995.

[13] M. Nielsen and G. Winskel. Trace structures and other models for concurrency. In V. Diekert and G. Rozenberg, editors, *The Book of Traces*, pages 271–305. World Scientific, Singapore, 1995.

[14] D. Peled. Partial order reduction : model checking using representatives. In *MFCS'9*, volume 1113 of *LNCS*, pages 93–112, 1996.

[15] D. Peled, T. Wilke and P. Wolper. An Algorithmic Approach for Checking Closure Properties of ω-Regular Languages. In *CONCUR'96*, volume 1119 of *LNCS*, Springer-Verlag, (1996) 596-610.

[16] A. Pnueli. The temporal logic of programs. In *18th Symposium on Foundations of Computer Science*, pages 46–57, 1977.

[17] R. Ramanujam. Locally linear time temporal logic. In *LICS '96*, pages 118–128, 1996.

[18] P. S. Thiagarajan. A trace based extension of linear time temporal logic. In *LICS '94*, pages 438–447, 1994.

[19] P. S. Thiagarajan. A trace consistent subset of PTL. In *CONCUR '95*, volume 962 of *LNCS*, Springer-Verlag, (1995) 438-452.

[20] W. Thomas. Automata on infinite objects. In J. van Leeuven, editor, *Handbook of Theoretical Computer Science Vol.B*, pages 133–192. Elsevier, 1990.

[21] W. Thomas. On logical definability of trace languages. In V. Diekert, editor, *Workshop of the ESPRIT Basic Research Action No: 3166*, volume Report TUM-19002, Technical University of Munich,, pages 172–182, 1990.

[22] A. Valmari. A stubborn attack on state explosion. *Formal Methods in System Design*, 1:297–322, 1992.

[23] L. Zuck. *Past temporal logic*. PhD thesis, Weizmann Institute of Science, Israel, 1986.